NASA-CR-197973

/NTERIM
7N-55-CR
OCIT
44787
P. 21

Wide-Bandwidth High-Resolution Search for Extraterrestrial Intelligence

Status Report

15 Dec 1993 -- 15 Dec 1994

Paul Horowitz, Principal Investigator

Harvard University
Cambridge, MA
02138

January 15, 1995

Grant Number NAGW-2872

(NASA-CR-197973) WIDE-BANDWIDTH HIGH-RESOLUTION SEARCH FOR EXTRATERRESTRIAL INTELLIGENCE Status Report, 15 Dec. 1993 - 15 Dec. 1994 (Harvard Univ.) 21 p

N95-70962

Unclas

Z9/55 0044787

1. INTRODUCTION

This report summarizes research accomplished during the third year of the grant. During the period covered by this report the active personnel included the PI, two Harvard graduate students (Darren Leigh and Jonathan Weintroub) and an MIT graduate student (Max Avruch), each averaging a 75% research commitment; three Harvard undergraduates (Derrick Bass, Suhail Shah, and Eric Wey; the latter two during the summer only), and a recent mathematics graduate from Harvard (Nick Shectman).

2. RESEARCH ACCOMPLISHED

2.1 Final Experiment Architecture

As the construction of this SETI project ("BETA") has evolved, and based on our growing experience with the data from the ongoing project META, the system architecture has been evolving, generally in the direction of greater robustness against terrestrial interference. Thus, for example, we added parity checking throughout, not only for all DRAM operations, but also for every signal-bearing cable. We expanded the original 2-beam design to a 3-antenna configuration, adding a terrestrial low-gain discone feed that can "veto" an otherwise interesting spectral feature.

We have also increased the channel count -- from an original objective of 100 million channels to 240 million (plus 12 million redundant "rover" channels). Finally, we have greatly enhanced the backend "feature recognizer": Originally it was a simple, and inflexible, baseline/threshold circuit, with the additional capability of following interesting frequency regions ("slots"), and ignoring others ("notches"). The new design, almost complete, adds a powerful state-memory feature, multiple simultaneous thresholds, and the ability to integrate multiple spectra in a flexible state-machine architecture. This is described in detail below.

Figure 1 is the block diagram of the BETA architecture; we intend to complete construction and put the system on-line this year. Thus the architecture is now "final."

2.2 FFT Array

In previous reports we have described our innovative 3-chip design for a 4 megapoint complex FFT, wherein the long transform is implemented as a succession of shorter row and column FFTs, with complex ("twiddle") multiplies interposed between the shorter transforms during the "corner turns" (an evolution of the elegant Serendip design of the Berkeley SETI group). We also described numerical simulations, verifying detection of weak narrowband signals in the presence of much stronger signals and broadband noise (see the 12/92 progress report for an exhaustive description). Those simulations established the degree of window- and twiddle-ROM truncation that can be used (to reduce cost), consistent with desired dynamic range and spur-free response.

2.2.1 FFT Hardware Verification

In the last progress report we described the implementation of the FFT architecture as finished working hardware, and the exhaustive testing of the prototype PC board with various composite signals. These included weak sinusoids embedded in noise, closely-spaced tones of large amplitude ratios, and modulated carriers; in all cases we tested the performance in the most difficult cases, namely "mid-bin" test frequencies (i.e., signal frequencies incommensurate with the overall sampling window), in order to evaluate the contribution of windowing to spectral leakage in the worst case. For example, those tests demonstrated accurate spectral display of a weak mid-bin carrier separated just 10Hz (20 channels) from a mid-bin carrier 40dB stronger, both in the presence of additive wideband noise.

The hardware includes several diagnostic aids, namely 1) a set of headers for monitoring the passage of data through the FFT, along with a deterministic hardware data simulator for which we tabulated "signatures" at each header, and 2) a set of addressable test ports at input and output, so that the spectrum of a pair of FFT boards analyzing the identical input data stream can be compared for precise bitwise equality of output spectra. We have verified (with a logic analyzer) that the observed data at all the headers matches the precalculated signatures; the use of IC sockets throughout should make maintenance straightforward.

2.2.2 FFT Array Production

During the period of this report we completed production of the full 63-board array (plus 7 spares), thanks especially to gifts from Micron Technology (of the required 3000 MBytes of DRAM) and from AMD (of Mach CPLDs). Skilled and dedicated student solderers loaded and soldered the quarter-million solder joints, then defluxed the boards and stuffed the components. We constructed card cages, power supplies, and all interconnecting cabling. At initial startup 90% of the boards worked perfectly. With previous completion of the LO Array, IF Channelizer, and Mixer/Digitizer Array, and with completion of the power supplies and card enclosures, BETA's 250 Megachannel Spectrometer is now complete! See Figure 2, which shows the entire spectrometer (except for cabling) in its single rack.

2.2.3 FFT Control

The spectrometer array is strictly synchronous, with 40 MHz differential ECL clock distribution generated from the 10 MHz GPS-derived master station oscillator (from which all the LO synthesizers, as well as the backend Pentium clock signals, are also derived; this forces all oscillator feedthrough -- observatory "birdies" -- to be aliased to DC in every spectrometer board). Synchronization signals to the individual FFT boards are similarly distributed, ensuring that all FFT processors proceed in lock-step. Figure 3 shows the block diagram of the master FFT control, fabricated on a ground-plane wirewrap panel.

The FFT control also performs the pairwise board validation, as can be see in Figure 3: A pseudorandom test vector generator (implemented on a Mach 210 CPLD) drives a selectable pair of FFTs, each component of whose outputs are compared for equality. The sequencing of board pairs, and the verification of equality, are controlled by an ISA interface on the system's master PC.

The control hardware, including the CPLD code, is complete. We expect to debug this essential system shortly.

2.3 Feature Extractor

The photograph in Figure 2 contains all the hardware upstream of the blocks labelled "Feature Recognizer" etc., in Figure 1. This is a crucial component of the entire system, given that the unexpurgated output of the spectrometer that feeds it consists of 250 MByte/sec of spectral amplitudes. No general purpose commercial hardware can handle data at this rate. The Feature Extractor is responsible for maintaining a moving baseline, recognizing large spectral peaks, following the progress of previously identified interesting spectral regions, and blocking signals from regions previously identified as containing interference. In its current incarnation it can also correlate nearby in frequency and time, and (in integration mode) it can average many spectra in a tracking reobservation.

The hardware configuration consists of 21 copies of three Feature Recognizers (FR), one Feature Correlator (FC), and one memory board, each set of 5 boards resident on one Pentium ISA bus. This 5-board set receives the spectra from three FFT boards -- fed from the three antenna feeds, and covering a common 2 MHz portion of spectrum. Thus there are 21 such sets (20 of them cover 40 MHz of instantaneous contiguous spectrum, and the 21st is a redundant "rover" set).

Figure 4 shows the block diagram of a set of 5 boards on its Pentium motherboard; in Figure 5 we show a photograph of the three kinds of boards, in this instance undergoing serious debugging.

In brief, the function of the board set can be described as follows: The Feature Recognizers associate a baseline average and a set of four threshold comparison results (against the baseline) with each data point. The Feature Correlator considers the comparison results from each channel in a 2-dimensional (time and frequency) state machine and decides whether to report (to the PC) data (and associated baselines) from the Feature Correlators. In addition, the memory board allows one of the FRs to perform integrations and two modes of readout: a direct readout of memory to the PC, and a "recirculated" readout which applies the feature-detection algorithms to the integrated data. All systems are easily configurable from the PC, and the FC can be directed to consider certain areas of the spectrum as more or less interesting.

The Feature Recognizer and memory board have been fully debugged and are being built in production quantities. The Feature Correlator has been built in prototype quantities and is in the final stage of debugging.

The Feature Extractor operation is the most complex portion of the entire system. We have included a detailed description as an Appendix to this report.

2.4 Pentium Array and Workstation

The Pentium array consists of 21 Pentium-based PC motherboards, each with 16 MByte of RAM and an Ethernet interface. The motherboards have neither disks (hard drives or floppies), monitors, nor keyboards. Each receives and processes the data from a feature extractor/correlator board set, passing on the results of a first analysis to the central Unix workstation (through which each is also booted). The array's code is written in C++ using a DOS extender, thus providing a "flat" memory model (no segments or 640K barrier). Since high speed and low latency are important requirements, the array software deals intimately with the Pentium motherboard interrupts. The C++ programs are linked with hand crafted assembly language subroutines which handle these. The result is a fast, clean interface between the high level programming code and the motherboard hardware. This code has been exhaustively tested via communication between a PC and the workstation.

Thanks to gifts of DRAM (from Micron Technology Inc.) and Pentium microprocessors (from Intel), the back-end array is now up, running and communicating across the network. Each motherboard is housed in a separate mini-tower case (with integral power supply) and linked into the whole system by thin-wire ethernet (Figure 6). These diskless computers boot up via the network, treating a file on the Unix workstation as a 1.44 MByte floppy disk. To increase the versatility of the computer array, individual computers can be booted from different "virtual" floppy disks, allowing them to run programs specific to an individual task. This capability is especially useful for programming the "rover" PC.

To ease software development for these PCs, we have supplied the workstation with modified utilities for accessing and manipulating MS-DOS format disks (and their virtual images). Using these, the programmer can easily copy programs from a PC to the network accessible virtual disks, set up test procedures and run them. To the running software, a back-end computer appears identical to a computer with a floppy drive.

The real-time PC is now mostly complete. It talks to the GPS station clock and the first LO's frequency synthesizer (in order to compensate site Doppler effect) and sends synchronization packets across the ethernet. We have developed celestial modeling software to compensate doppler effect to the needed accuracy. This software takes into account the earth's rotation and oblateness, barycentric motion relative to the earth-moon system, and motion around the heliocenter. This results in doppler offset calculations to 5 m/s accuracy over the long term. Over the short term (several beam durations or integration periods of a few minutes) the effect of residual doppler errors are negligible, corresponding to less than a single spectral channel smearing. Display software for this PC shows useful real-time quantities such as the current UTC and sidereal times, antenna pointing, and current frequency band under observation.

Low level workstation control software is now complete and nearly fully debugged. The workstation is capable of receiving ethernet packets from the real-time PC to synchronize the system and obtain information about antenna position and real time parameters. Upon receipt of the real-time PC packet, the workstation interrogates each Pentium in the backend array, receiving data about new hits received and slots that have been completed. There are also hooks to allow the workstation to request premature slot downloads and frequency notching. By dynamically monitoring the quantity of data received, it can modify threshold values for each Pentium separately to maintain optimum network and CPU loading.

Because we do not yet have data flowing through the newly-assembled Pentium array, the "real-time Pentium" and the "backend array" are simulation programs residing on the host workstation. The backend array simulator generates signals based on the statistics of gaussian noise with embedded monochromatic sidereally scanning sources (i.e. proposed ETI signals).

Currently we are investigating data analysis and archival algorithms. Preliminary research suggests that in low wind situations (the radio dish can be blown by approximately one beam width under wind stress) where scintillation is negligible over a dual beam traversal, a fitting of received data to the true dual beam profile shows promise for detection of sidereally scanning signals. The computational expense of the method as well as its applicability to scintillating sources or during less than optimal weather conditions have not yet been investigated. Since the dish is fully steerable and devoted to this experiment, a fast method of distinguishing especially meritorious candidate signals will allow rapid revisiting (within minutes) of these signals.

A flexible means of data archival is especially necessary. Correlation between low level signals seen at the same location in the sky may warrant a closer look, while repeated signals at the same frequency but different spatial locations indicates recurrent terrestrial noise. Other ways of correlating data may also be important, so above all we would like the database to admit possible analysis forms that have yet to be thought of. We are currently investigating the flexibility and efficiency of several relational database packages.

2.5 Radioastronomy Spectrometer

In the previous progress reports we described the design of a combined spectrometer and power-accumulator, for use at Arecibo Observatory to search for neutral hydrogen emission from condensations of neutral hydrogen at high redshift (z=5). This project, a technological spinoff from our SETI work, is progressing well. Our proposal to use the Arecibo radiotelescope for the search has been approved by NAIC. The work is supported by non-NASA funding, and is a collaboration with Bernie Burke of MIT, Mike Davis of NAIC, and Jim Cordes of Cornell (who will be using the spectrometer for pulsar studies and also in the Cornell undergraduate astronomy program.

Ian Avruch, Darren Leigh, and Jonathan Weintroub traveled to the Arecibo Observatory in January 94. Dual helical feeds had been mounted on the catwalk before our arrival (Figure 7), and we installed a scaled down version of our system, based on four power spectrometer boards. Unfortunately, significant interference was encountered in the band of interest, including strong narrowband RFI from television stations and a perplexing periodic broadband source present, not only in our receiver, but also at other frequencies. This interference turned out to be due to a defective power supply in a laser system on the observatory site. The RFI made for difficult analysis of the several nights of data gathered. We could not unequivocally identify a true celestial source, although we did notice the system temperature rise as the beam swept through the Galactic plane. In view of the generally poor quality of the data, we elected to concentrate on making systematic improvements, rather then elaborate data analysis to try and extract faint sources.

On returning to Cambridge, we revised the spectrometer circuit board layout, and fabricated a production run of sixty boards. These have been assembled and tested, and a subset of 32 of these form the spectrometer core for the experiment. The balance are used for spares, Jim Cordes' pulsar spectrometer, and the Cornell undergraduate program. The Local Oscillator Array and the IF Channelizer (both instruments originally designed for SETI) were also built during the Spring of 1994. The electronics department at Arecibo assembled and tested the sixteen mixer-digitizer boards required for the experiment. In addition we have put together an analog front end for the system which allows us to operate completely independently of shared Arecibo electronics. Dan Werthimer of Berkeley kindly gave us a design for an IRIG interface board (to interface our PC to the observatory time standard); we built and tested two of these.

We returned to the Observatory in June 1994 and during the following four weeks completed the installation of the full-blown "back end" system (Figure 8). The periodic broadband interference so troublesome in January had been eliminated at its source, while a particularly troublesome narrowband television carrier at 212 MHz had been mitigated (through judicious retuning of the front end bandpass filter). The spectrometer and backend computer system works well: The system control and data logging software has been substantially improved, with data display and automatic test features added. The program runs on an IBM compatible PC. Data is logged to the hard disk drive, then automatically spooled daily to 2GByte digital audio tape (DAT); the physical tapes are shipped periodically to Cambridge (as the overall data rate is too high for an Internet transfer to be practical).

The front end is still too noisy. This is due primarily to ground spillover, in addition to lossy elements (bandpass filters and cables) before the first low noise gain stage. The ground spillover will be reduced through redesign and repositioning of the feeds. We are presently building a scale model (at 600MHz) of a new 8-element broadside dual polarization dipole array (after the suggestions of Mike Davis and Bernie Burke). ELNEC models and our own analysis of the feed's interaction with the spherical reflector have predicted that this design should achieve the desired gain and noise figure. We will verify the ELNEC predictions by direct calibration of the scale model, and make modifications if necessary, before building the full size unit. This we hope to have ready by March 1995. We continue to make other improvements, most recently installing new bandpass filters of very low loss in the receiver front end.

Jim Cordes (Cornell/NAIC) has built and tested one spectrometer board. It has been used in an undergraduate course at Cornell to make observations of Galactic hydrogen. A sample of results are presented as Figure 9. We have built twenty-four spectrometer boards for his pulsar spectrometer and undergraduate efforts. He is presently building a high speed computer interface and an LO Array for the pulsar spectrometer.

3. NEXT STEPS

We have made very good progress during this reporting period, most significantly on the completion of the 250 million channel spectrometer (and its spinoff radioastronomy spectrometer), the nearly complete debugging of the complex Feature Extractor system, and the assembly of the Pentium backend array. Corporate gifts, most notably from Micron Technology, have prevented our budget shortfall from blocking completion of the full BETA system.

The major remaining tasks are the production of the Feature Extractor boards, and their integration and debugging in the Pentium array. We intend to finish this year, and move the completed system out to the radiotelescope at Agassiz Station.

4. OTHER FUNDING

During the period of this report we have received funding from The Planetary Society, and the Bosack/Kruger Charitable Foundation, in addition to our grant of partial support from NASA. The Radioastronomy projects are separate from the NASA activities, with support from the two private sources above, and additional funding from NSF and NAIC.

5. PUBLICATIONS AND TALKS

Paul Horowitz gave talks on SETI at MIT's Media Lab, and at the Symposium in honor of Carl Sagan's 60th birthday at Cornell. Jonathan Weintroub gave a talk on the search for high redshift hydrogen at the NAIC Symposium II at Cornell; he and Max Avruch spoke on the same subject at the Arecibo summer program.

We continue to enjoy media coverage, variously in newspapers, magazines, radio, and television. Most recently we were filmed for a Scientific American Frontiers program (hosted by Alan Alda), which aired in the fall of 1994 (along with a half-hour SETI feature on a PBS series called "Future Quest").

6. ACKNOWLEDGEMENTS

We are most indebted to Micron Technology for their generous donation of approximately 3000 MB of DRAM SIMMs, which makes completion of the project possible. Intel generously donated Pentium processors, and AMD came through with the needed CPLDs. These gifts, together with equipment donations from Fluke and Hewlett Packard, have boosted morale and kept our project healthy.

For the high redshift search we are grateful to Edgar Castro and the electronics department at Arecibo for their significant efforts, including building and installing the helical feeds, providing much of the RF and IF analog electronics, and funding and assembling the mixer-digitizer array. Darrel Emerson (NRAO) provided data from ELNEC models of the helical feeds. He and Bob Zimmerman (NAIC) has provided a great deal of advice generally. In the area of computers and networks, the assistance of Arun Venkataramen (NAIC) has been invaluable.

APPENDIX

Project BETA -- Signal Recognition and Discrimination

OVERVIEW

See the accompanying block diagram in Figure 1 ("BETA Architecture"). BETA (Billion channel ExtraTerrestrial Assay) is a meridian transit search for narrowband (carrier) microwave beacons in the "waterhole" band of 1400-1720 MHz. It is an evolution of the currently-operating META search (8 million channels, 0.05 Hz RBW) to a quarter billion channels and three antenna beams. It uses an 84-foot Cassegrain radiotelescope with dual feedhorns (and a third low-gain terrestrial discone) feeding HEMT low-noise amplifiers in an agile heterodyne receiver with a 240 million channel Fourier spectrometer (80 million channels of 0.5 Hz RBW [i.e., 40 MHz instantaneous BW] for each feed), whose outputs (power spectra) feed an array of "feature recognizers" and "feature correlators." The latter sift through the 250 MByte/sec (approximately) of spectral data, seeking distinctive spectral features that may come from extraterrestrial beacon transmissions. Such signal candidates should transit from the East to the West horn (each has a half degree beamwidth, spaced about 1 degree; at equatorial declination the beamwidth corresponds to 2 minutes of time), without appearing in the low-gain discone.

BETA searches 320 MHz, as 8 "hops" of 40 MHz (the system's instantaneous BW). Each hop takes 2 seconds, thus a full cycle through the waterhole takes 16 seconds. With a beam transit time of about 2 minutes, each potential source is revisited about 8 times at each frequency hop, in each sky beam. If a particularly good candidate is seen (appears first in the East horn, then in the West, and then disappears, never being seen in the discone), we can break off the transit survey, "leapfrog" the antenna a few beamwidths to the West, and ask the source to perform an encore.

The BETA backend was designed so that we can look for good candidates, with rapid reobservation, while rejecting interfering signals. Examples of the latter are signals that appear simultaneously in both horns, or in a sky horn and the discone, or in a sky horn for longer than the transit time. It has the ability to look for strong spectral features ("hits"), follow interesting frequencies (a "slot") that were initially triggered by a hit, and to reject frequencies that are believed to contain persistent interfering signals (a "notch"). It can operate in several modes, including a state memory mode, and an integrating mode (for one hop band only, with full or partial spectral readout). Here follows some details of the planned backend operation.

DATA PATHS

Each Fourier transform feeds its data (4 million magnitude samples, 16 bit precision, in 2 seconds) into a "Feature Recognizer" (FR). This device recognizes peaks in the power spectrum. FFT boards and their associated FRs are arranged in groups of three, corresponding to the east, west, and terrestrial antenna horns. Each triplet of FRs is connected to a "Feature Correlator" (FC) which determines whether the data should be forwarded to its backend PC. The PC's perform low-level programming of the FCs and gather data from multiple passes through a given frequency window (hop). They forward interesting data across a LAN to a workstation, which archives and analyzes data in a higher level manner.

There are 20 such paths (i.e. 60 FFTs, 60 FRs, 20 FCs, and 20 PCs) operating in parallel, each covering a 2 MHz bandwidth, for a total instantaneous bandwidth of 40 MHz at 0.5 Hz resolution. There is one additional path which operates at the same frequency with one of the other 20, so that errors may be detected.

FEATURE RECOGNIZER

A 4K-point running baseline is computed by summing the squares of the magnitude data points received from the FFT. The data point at the center of this "moving baseline" (boxcar averaged) is forwarded along with the baseline value. This power sum is barrel shifted, converted back to magnitude, and multiplied by four programmed threshold values. (Actually, the multiplication table is stored in a ROM, so arbitrary 2 variable functions can be preprogrammed; the square root function is also encoded in a ROM, and can therefore be an arbitrary one variable function.) Both the amount of barrel shift and the value of the thresholds are programmable on the fly by the PC. Every element of the magnitude spectrum is compared with all four baseline-times-threshold values.

The largest threshold will represent a signal that is well above the noise baseline, while lower thresholds can be used to track signals of smaller intensity without necessarily forwarding them to the PC, as will be discussed in the Feature Correlator section.

The FRs also have an integration mode, in which power spectra can be summed and stored. This integrated power can then be read out directly to the PC, via the FCs. Because of the slow PC bus, an entire spectrum would require many seconds to fetch, and several spectra would be lost. The hardware therefore allows the PC to request only certain portions of the spectrum (the "slot" mechanism described below.) In addition, the integrated spectrum can be recirculated through the FR signal-recognition circuitry, just as an ordinary spectrum, allowing the same threshold analysis to be performed on the integrated spectrum; these partial readout methods allow quick readout from an integrated spectrum. The ability to read out the full spectrum would be useful in a sidereal tracking reobservation of many minutes' duration. Only one of each group of three FRs contains the RAM required to perform integrations.

Since the Feature Recognizers are independent units, they may be in separate modes. One interesting proposal has been to place the west horn in integration mode while the east is in normal mode so that interesting hits seen in the east horn can be then be integrated by the west to see if they are coherent. This mode of operation, however, would not allow the use of a sidereal scan filter, i.e. watching a signal handoff from the east horn to the west at a rate consistent with the sidereal scan of the antenna.

FEATURE CORRELATOR

The Feature Correlator forwards interesting frequencies to the PC. (The data it forwards are baseline and signal magnitude for the three horns as well as the frequency bin number.) The PC has the ability to set certain frequencies to be automatically forwarded (by specifying a start and stop frequency) -- "slots." When a good hit is seen, the PC may begin a slot, meaning that while the presumed source is still in one of the beams, the PC requests a frequency range around the original frequency be automatically forwarded to it. The inverse of a slot is a notch, a range from which the PC requests that no data be forwarded. Because of the limited bandwidth of the FC/PC interface, it may be important to notch known RFI to prevent loss of potentially interesting data due to an overflow of uninteresting noise.

If a frequency is neither a slot nor a notch, a programmable state machine (SM) determines whether it is interesting enough to forward to the PC anyway. The SM is a RAM whose program is configurable from the PC. Such a reprogramming will take the FC offline for 2 seconds. If, however, the FR is in integration mode, SM reprogramming can occur during integration, avoiding the need for any downtime.

The SM receives a 7-bit quantity describing how many thresholds were exceeded by the signal in the three lobes (4 thresholds => 5 states per FR, three of which gives 5^3=125 total states). It also receives a 5-bit quantity ("time based state") that it stored in the FR RAM the last time a particular frequency was visited (16 seconds previously), and a 2-bit quantity ("frequency based state") that it sent to itself from the previous frequency bin (0.5 microsecond previously). Based on this information, the SM generates 5 bits of time based state to be stored and retrieved the next time that frequency appears, 2 bits of frequency based state to send to the next frequency bin, and a bit specifying whether the data of this frequency bin should be forwarded to the PC.

The SM allows great flexibility in making intelligent decisions at speeds with which the PC cannot compete. For example, the frequency based state can be used to detect a signal that happens to be in between two bins, so that it does not exceed the highest threshold, but exceeds the second highest threshold for two consecutive frequency bins. The time based state can be used to monitor a signal that is crawling into the edge of beam, but has not yet arrived at the center and so has not exceeded a high threshold, but has repeatedly exceeded lower thresholds. The PC can be sent this frequency bin with the understanding that it may wish to follow it so that if it does turn out to eventually exceed threshold, less data will have been lost. At the moment, we do not know exactly what all four thresholds will be used for, nor do we know what all the states will mean. Much of the meaning may come from experience after the system is online.

Since the SM sees the comparison results from all three horns, it can perform rejection of terrestrial interference and of signals that are simultaneously high in both the east and the west lobes, which may indicate RFI that has somehow not been detected in the terrestrial (discone) antenna.

We have developed an intuitive language for generating SM programs in an easy to understand format.

PC ARRAY

Each FC talks with one Pentium based PC motherboard. The PC will tell the FCs which slots and notches they should track and will gather their data into a format acceptable to the workstation. Since the processing power of this backend is substantially greater than we had originally anticipated, we are hoping to offload some of the analysis to the PC's. It is likely that they will make local decisions, such as which hits should be turned into slots and followed. They may also perform signal analysis on the slots. However, because of the expense and complexity of maintaining more than 20 fully equipped PC systems, none of these PCs has a disk or monitor (they boot and load code via the LAN). Thus any data that must be archived needs to be forwarded across the LAN to the workstation. Such data includes actual slot downloads as well as information about possible sources of noise. The workstation will also receive planetary ephemerides from a special PC, the "real-time PC", and will be able to request notches so that we do not waste time repeatedly storing saturated spectra from the sun, for example.

The workstation may also perform analysis of slot data. At present it is unclear how much processing power will remain in the PC array after FC and network housekeeping jobs are taken into account.

Since the radio dish used in this project is dedicated to SETI and fully steerable, in a later phase of BETA it will be under computer control. Then, the workstation may direct a revisit of particular locations on the sky when a good candidate is found. This happens within minutes of initial signal reception. The FR's may be placed in integration mode during this time. In addition, the local oscillator array can be fixed (no hopping), allowing the signal level to be measured every two seconds rather than every 16. Revisitation will also provide another RFI rejection mechanism, since the relative placement of beam sidelobes with respect to the earth will change when the dish slews.

DATA ANALYSIS ALGORITHMS

Not much research has been performed in this area. We have a simulator that adds occasional ETI signals (modeled as one bin wide, possibly wandering signals) to a Gaussian noise base. Data fitting to the antenna beam pattern, along with a demand for sidereal transit between the east and west beams, seems to hold promise for signals that are not modulated in amplitude over a time scale on the order of several minutes (the time to transit both beams). We have not yet investigated the performance of this algorithm for sources that are scintillating with a time constant of roughly this order. A more serious problem is that our dish is somewhat "rubbery," and can wander a significant fraction of a beamwidth in moderate winds (we suspect this is why we have its continuous use for SETI!). We are considering absolute position sensors on the dish so that an attempt can be made to deconvolve signals obtained in such situations.

10/28/94

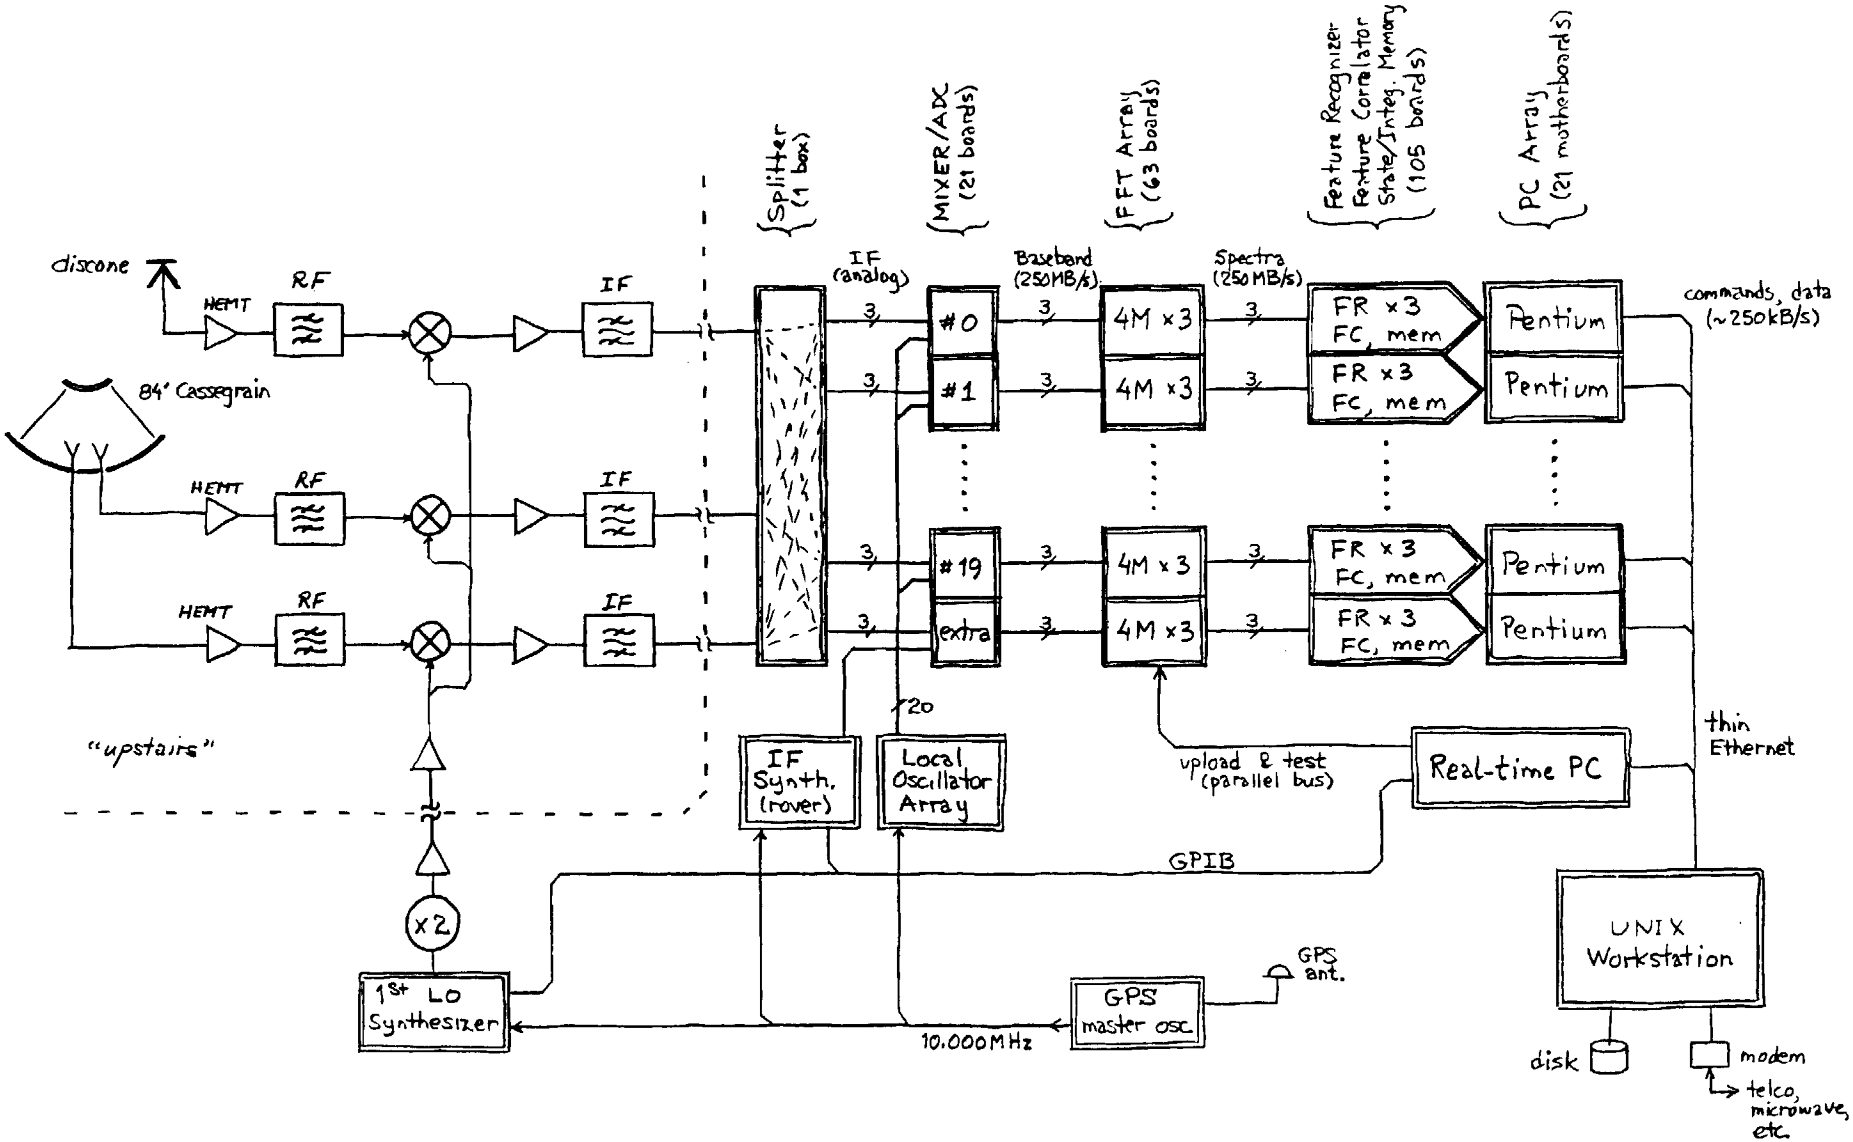

Figure 1. Final block diagram for BETA: Billion-channel Extraterrestrial Assay. The 240-million channel spectrometer allocates 80 million channels (40 MHz instantaneous bandwidth) each to a transit telescope with a pair of east/west sky lobes, and to a third low-gain feed (to veto interfering signals). The spectrometer's output (250 MByte/sec) is sifted by a flexible "feature extractor" array.

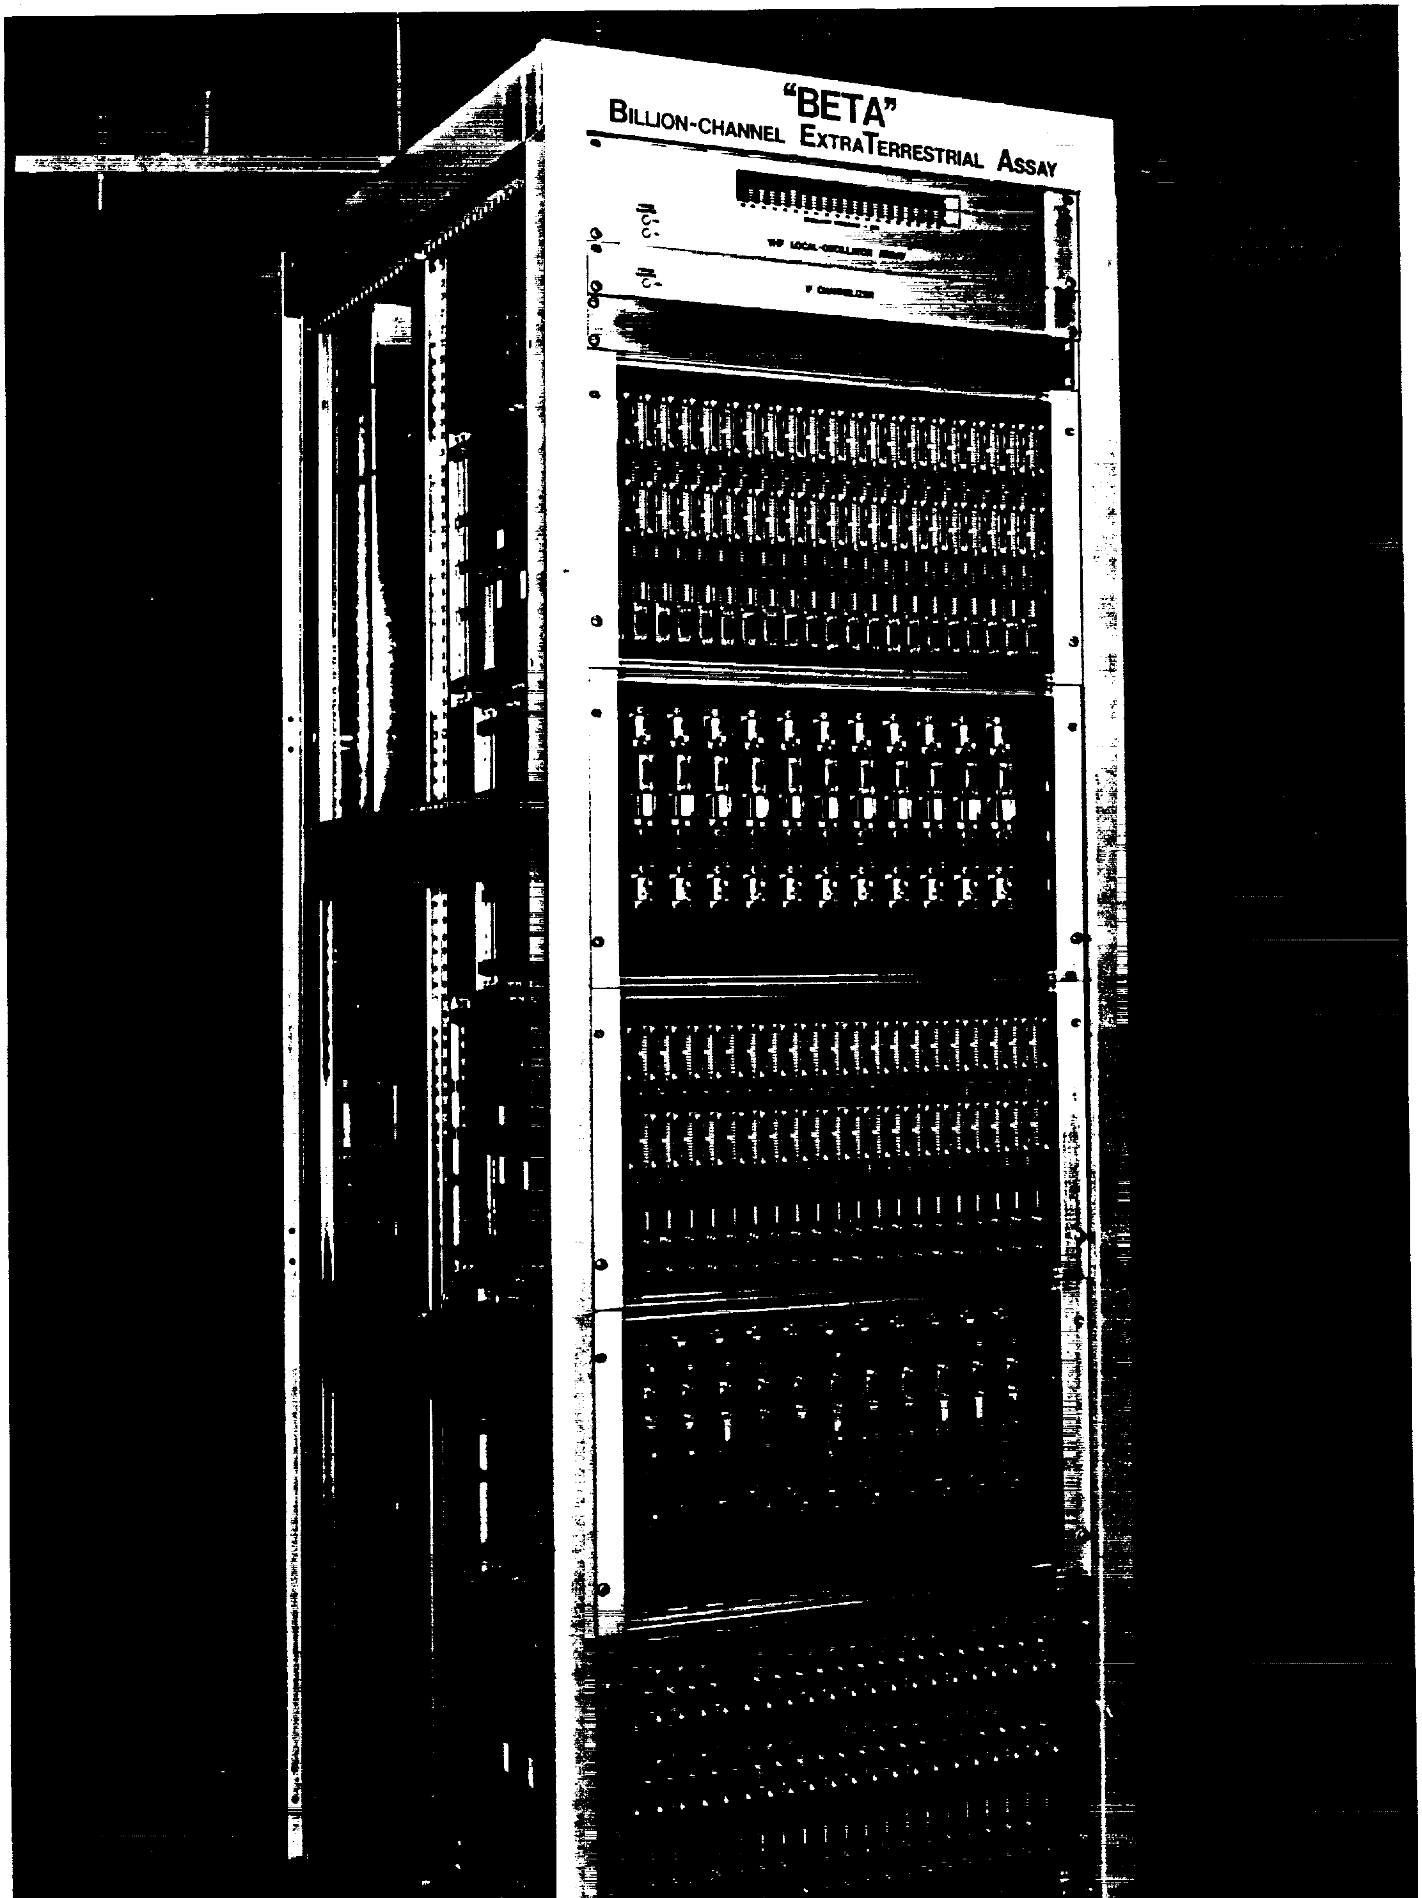

Figure 2. A quarter billion channel spectrometer (and upstream circuitry) in a rack. This array of 4 megapoint FFT boards performs a 240-million channel FFT every 2 seconds, corresponding to roughly 30,000 MIPS.

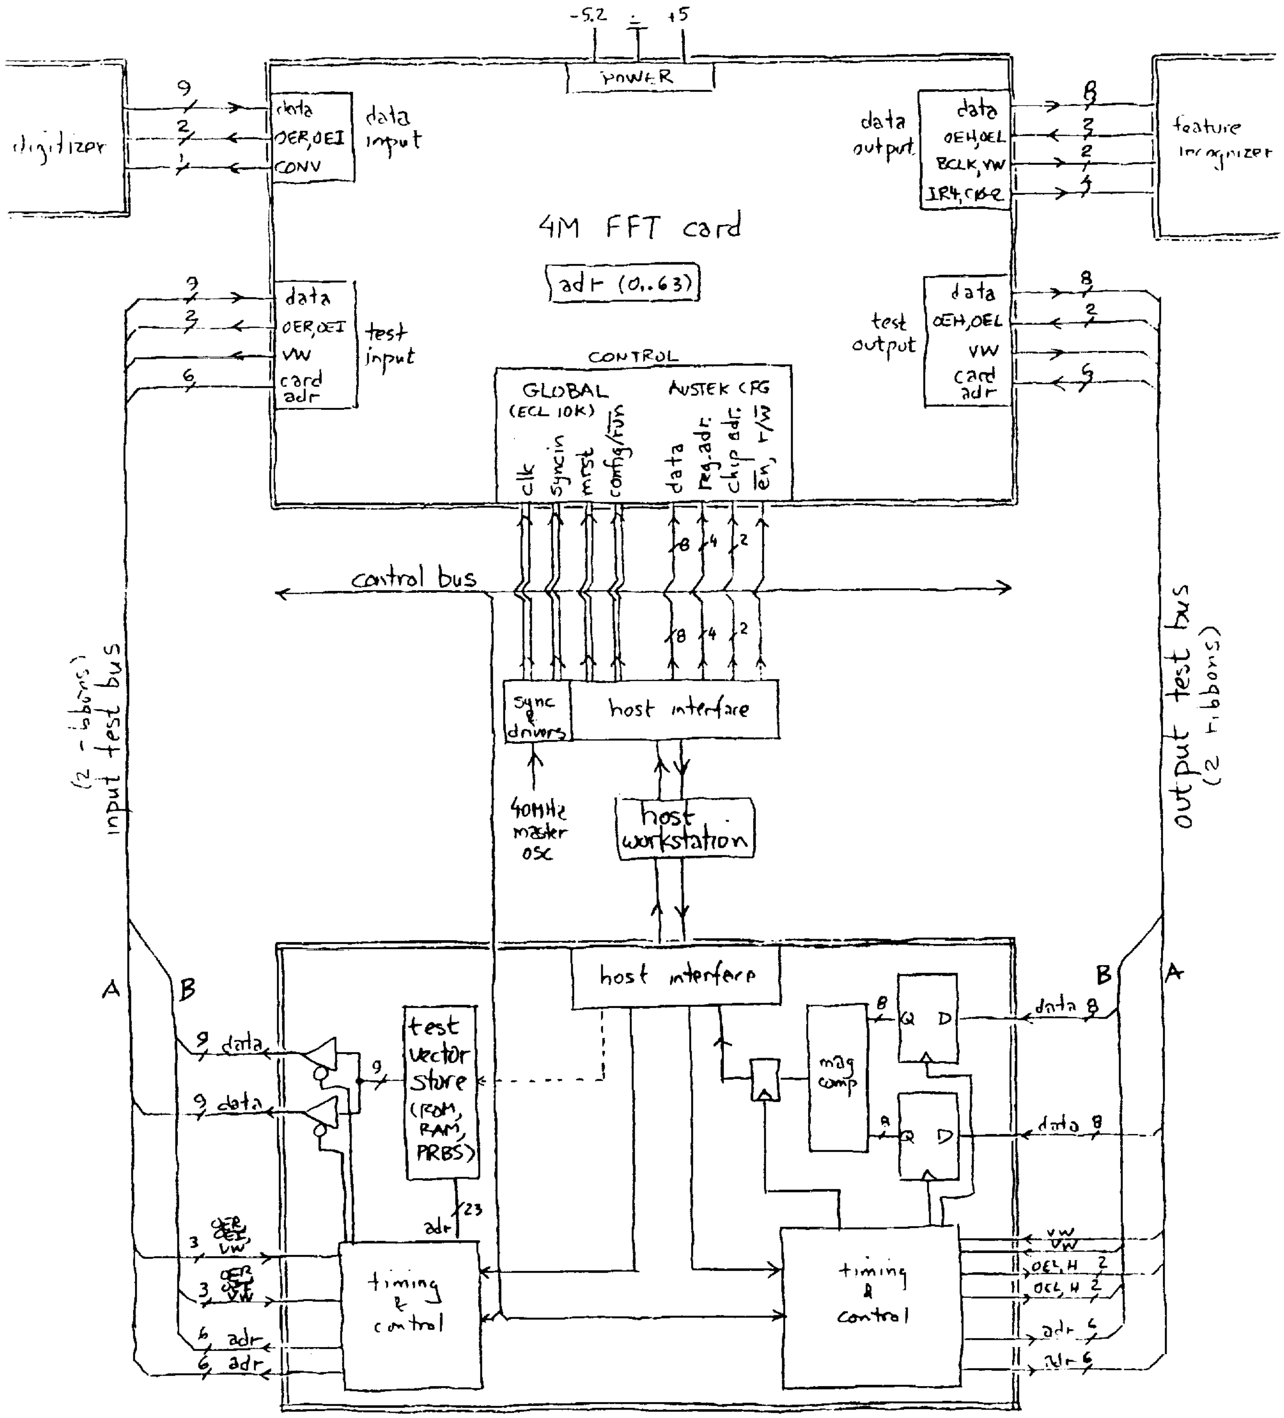

Notes:

1. 40MHz clk, & SYNCIN# "on-the-fly" distributed via ECL 10124/10125 bussed pairs; SYNCIN# carefully timed.
2. Austek upload from host; non-existent chip number is a register (20V8), to hold test/norm, clear, etc. state.
3. Austek upload & extra state bits set in "config"; transition to "run" lets PAL's begin counting.
4. All FFT bds are in lockstep; control-bus driver cktry "knows" C-states & BCLK timing.
5. In "test" mode, normal input/output disabled. Identical data is applied to a chosen pair (1 in A, 1 in B) of FFT bds, & the outputs compared (after latency — can be made "seamless")

Figure 3. Control and verification of the FFT array. Master distribution of clock and synchronization signals puts the array in lock-step. The addressable test ports are used for routine verification of FFT board function, demanding identical spectra for identical inputs.

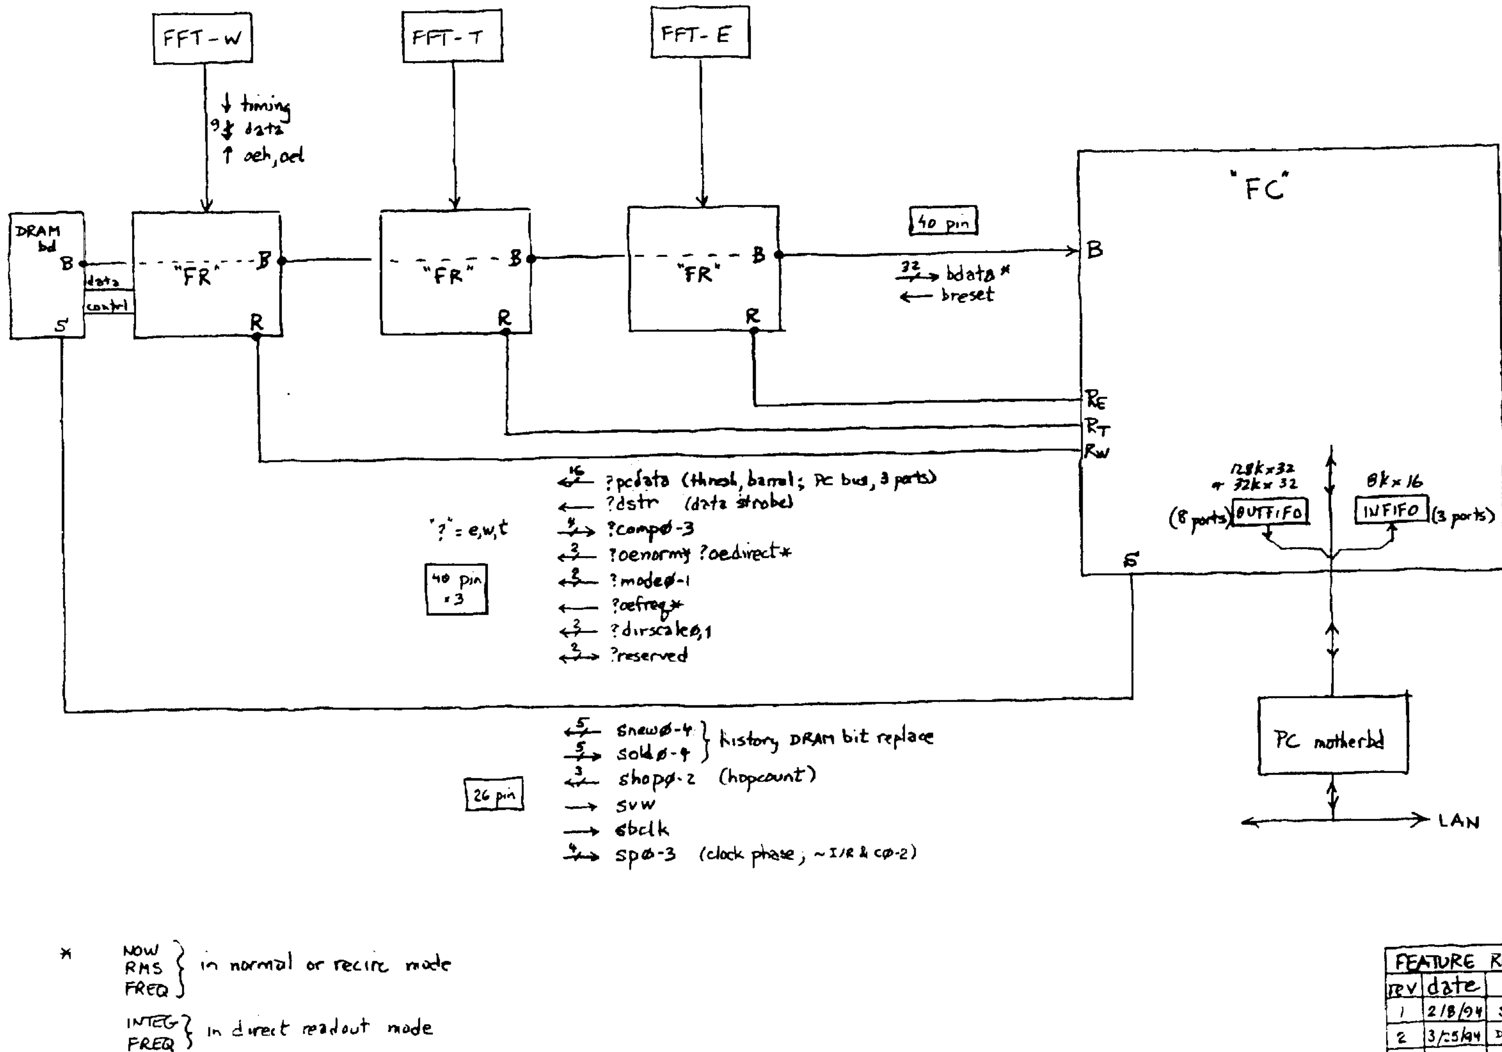

FEATURE RECOGNIZER/CORRELATOR		
rev	date	comments
1	2/18/94	signals named
2	3/25/94	DRAM bd, sigs cranked

Figure 4. Feature extractor arrangement: the three feature recognizers (FR) sift the spectra from a triplet of FFT boards, under control of a feature correlator (FC) that communicates with a Pentium-based PC. 20 MBytes of DRAM permits some memory of recent events (in time and frequency), or true spectral integration.

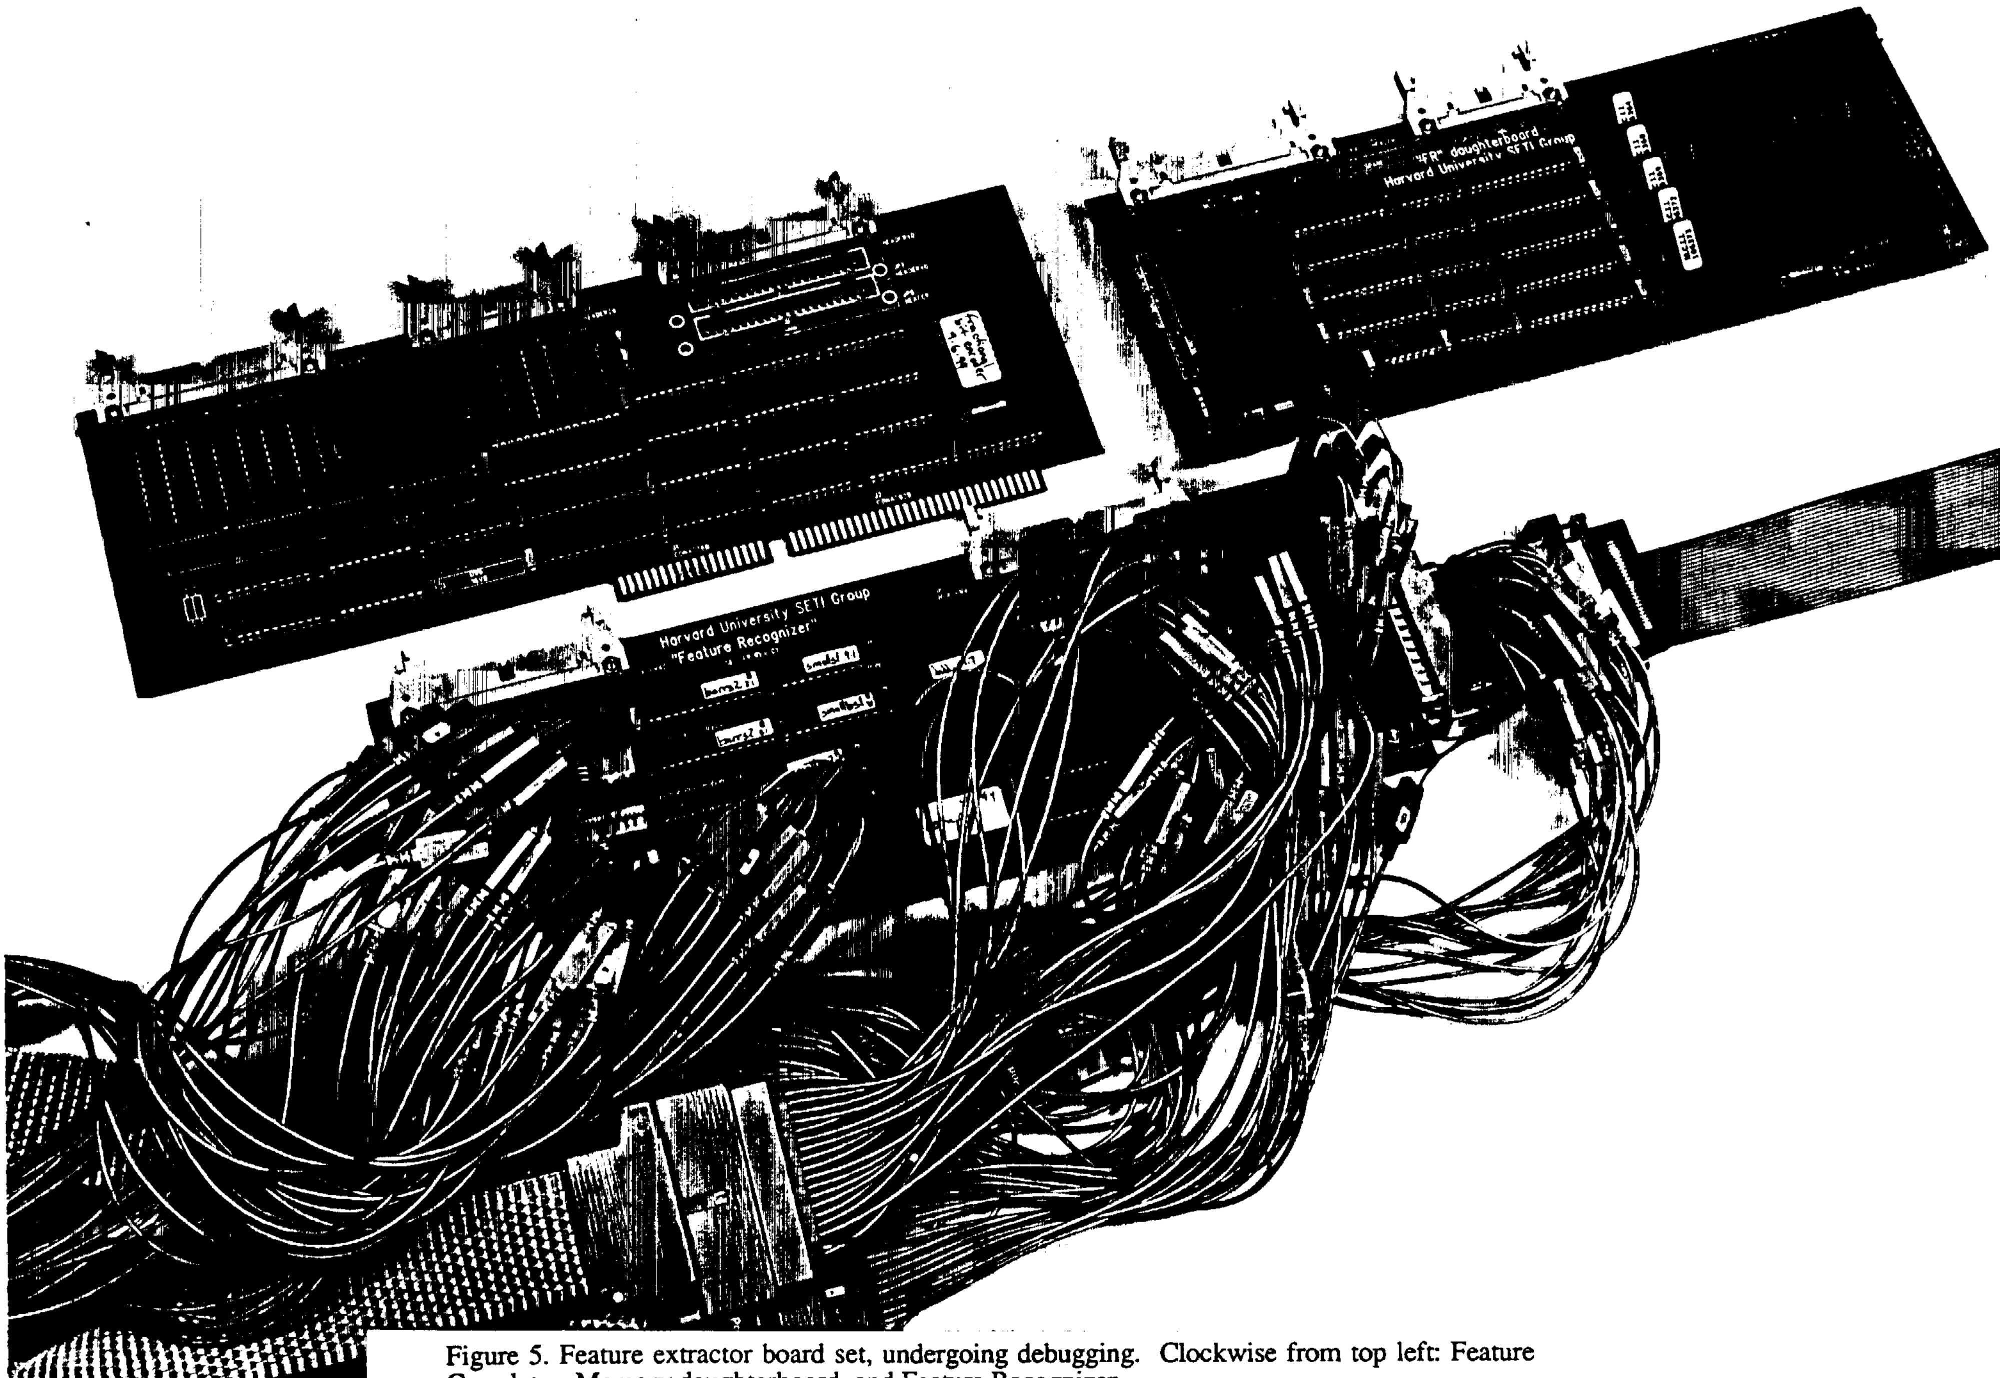

Figure 5. Feature extractor board set, undergoing debugging. Clockwise from top left: Feature Correlator, Memory daughterboard, and Feature Recognizer.

Figure 6. Array of Pentium-based Feature Extractors. Each mini-tower holds a 60 MHz Pentium motherboard, a 5-board feature extractor set, 16 MByte of DRAM, and a thin-wire Ethernet card. The diskless array boots from the LAN, receives data from the FFT spectrometer, carries out most of the decision making locally, and communicates with a UNIX workstation.

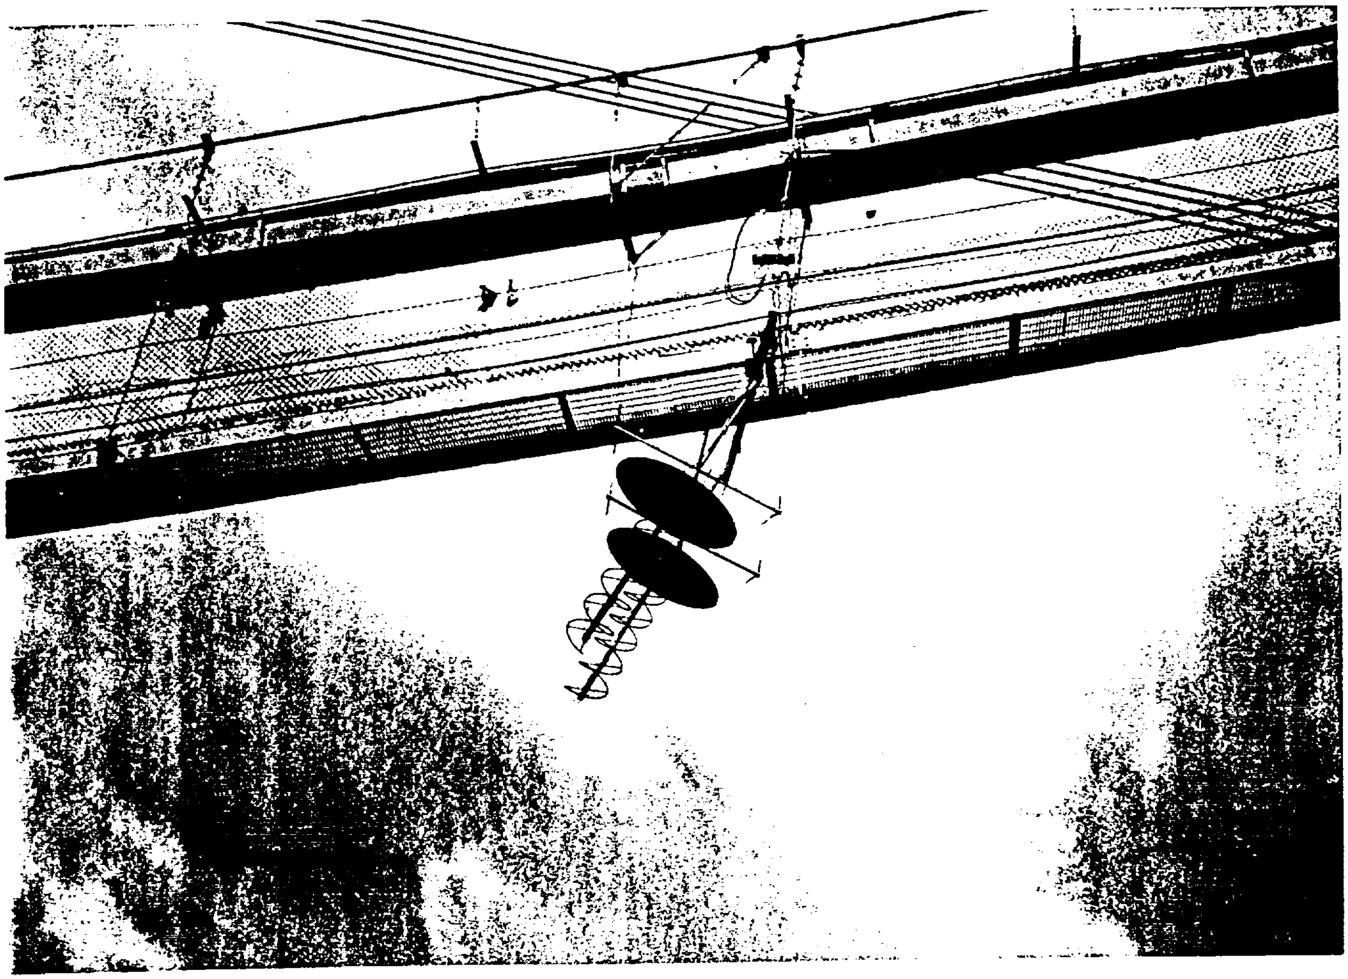

Figure 7. Dual helical feeds suspended from the catwalk at Arecibo Observatory, used in the search for high-redshift neutral hydrogen.

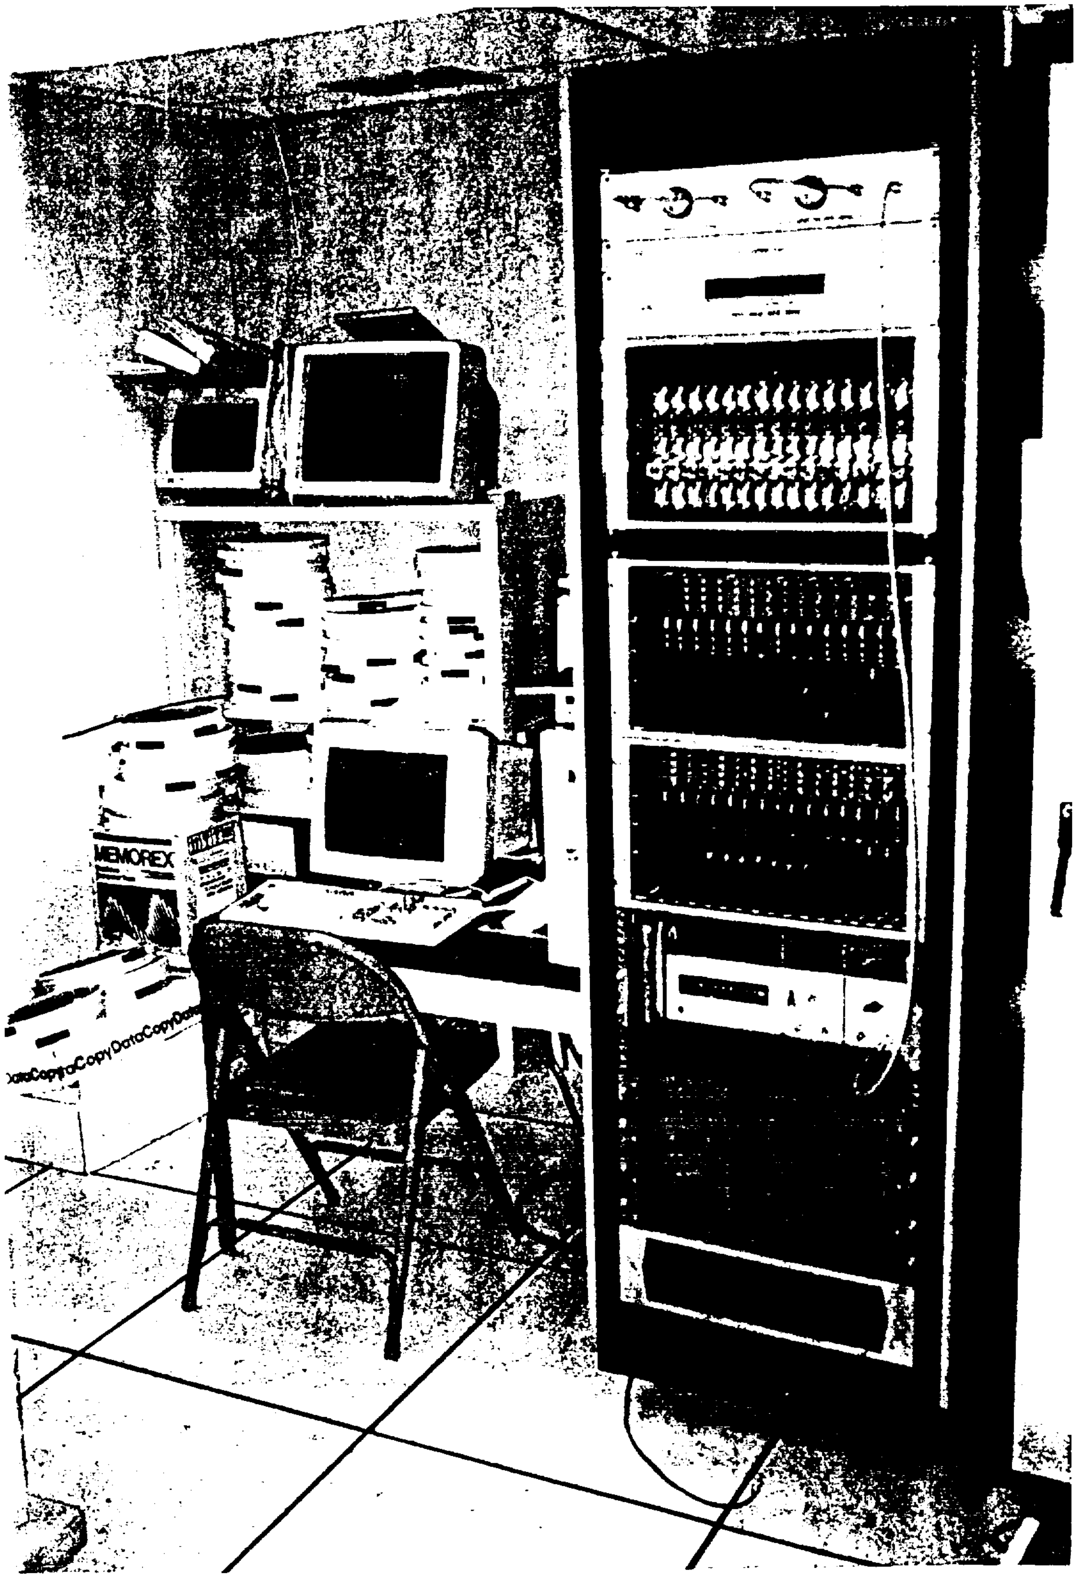

Figure 8. 8,192-channel dual-feed spectrometer at Arecibo. Instantaneous bandwidth is 32 MHz in each polarization. The rack also holds the local oscillators and quadrature downconverters. Data is archived on DAT, then shipped to Cambridge for analysis.

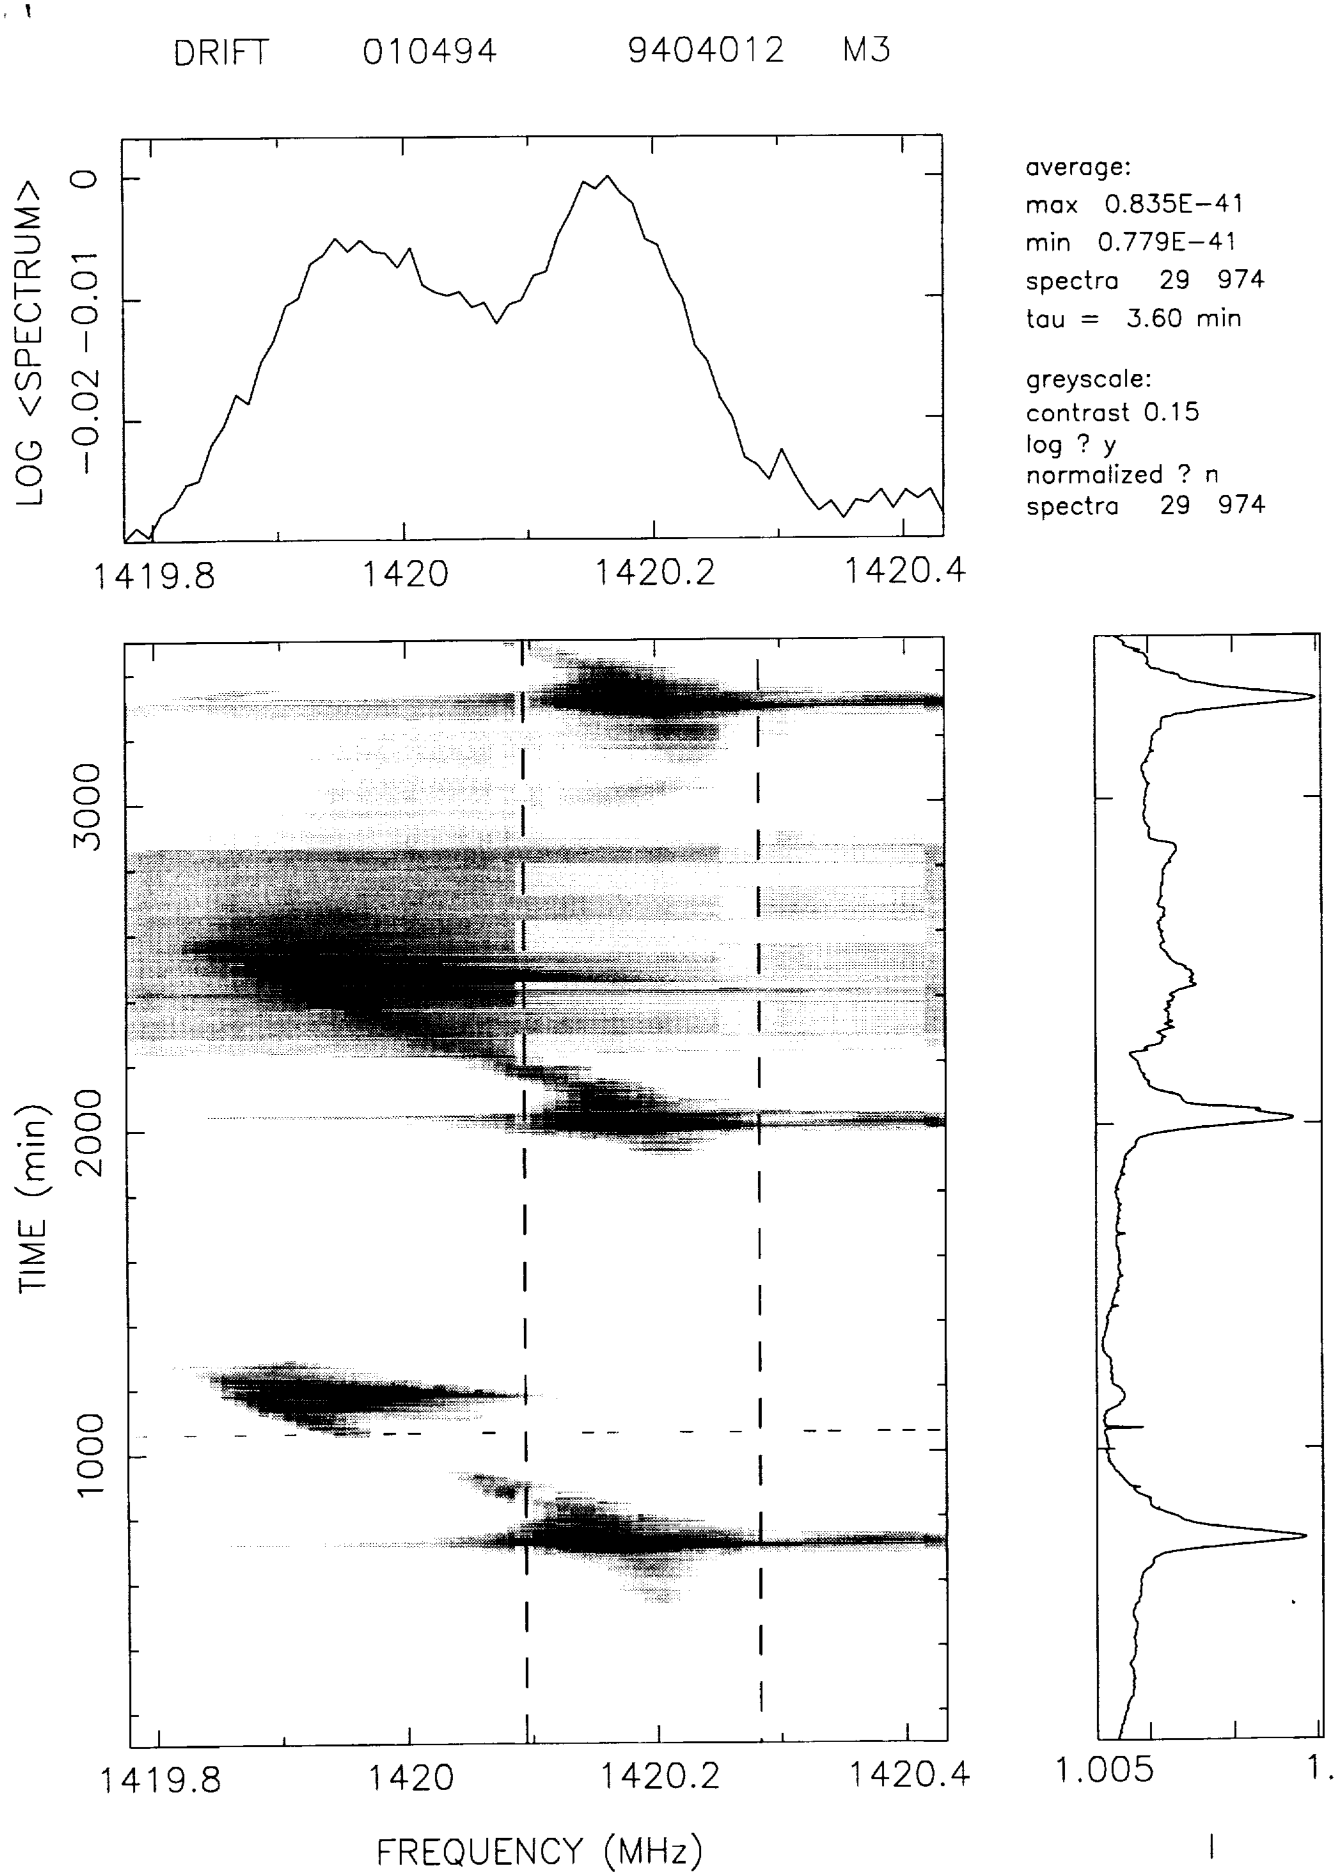

Figure 9. Hydrogen-line drift scan, from a Cornell rooftop (Cordes et al.), using one of the 256-channel integrating spectrometers built for the high redshift survey at Arecibo. The diurnal passage of the galactic arms is easily seen, as well as projections in the time and frequency axes.